30th International Summer Institute on Phenomenology of Elementary Particle Physics and Cosmology (SI2026)

Boltzmann Equation Solver for Thermalization

Jong-Hyun Yoon
Chungnam National University

Based on arXiv:2603.28848
Comput. Phys. Commun. 327 (2026) 110295
Code: github.com/best-hep/best

Fuji Calm, Fuji-Yoshida
12 August 2026

Sponsored by APCTP, Kavli IPMU, KMI and YNU

Outline

Part I — Reaching equilibrium

- Motivation: when do we need the full Boltzmann equation?
- Direct Monte Carlo for arbitrary n → m processes
- The identical-particle decomposition
- Validation & thermalization

Part II — Leaving equilibrium (preliminary)

- From flat spacetime to the expanding universe
- Sub-threshold freeze-out: relic abundance from the full f(p, t)

Summary

Particle production in the early universe

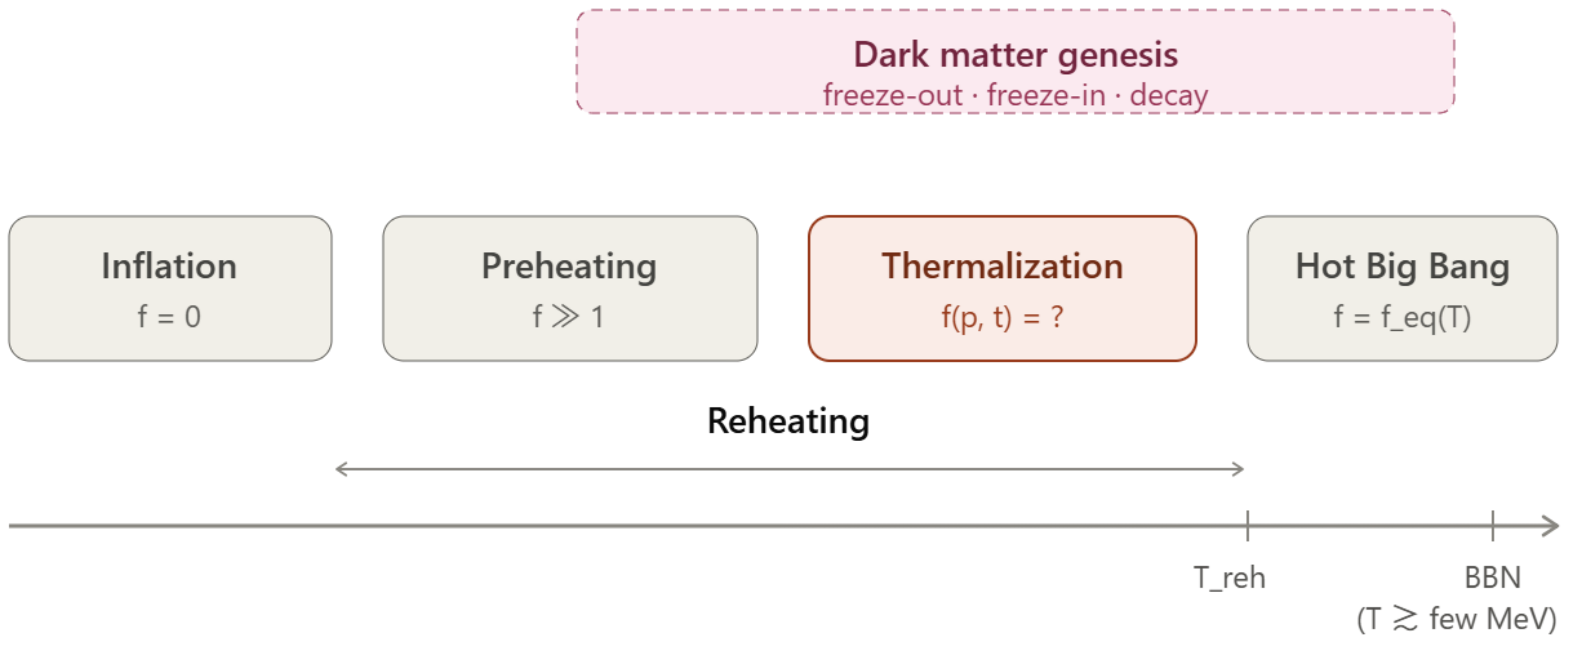

The reheating temperature

Unmeasured: anywhere between ~MeV (BBN) and ~10^{15} GeV and it decides which dark-matter and baryogenesis scenarios are even possible.

The standard estimate of the reheating temperature:

$$T_{\rm reh} \sim \left(\frac{90}{\pi^2 g_*}\right)^{1/4} \sqrt{\Gamma_\phi M_{\rm Pl}}$$

assumes instantaneous thermalization of the inflaton decay products.

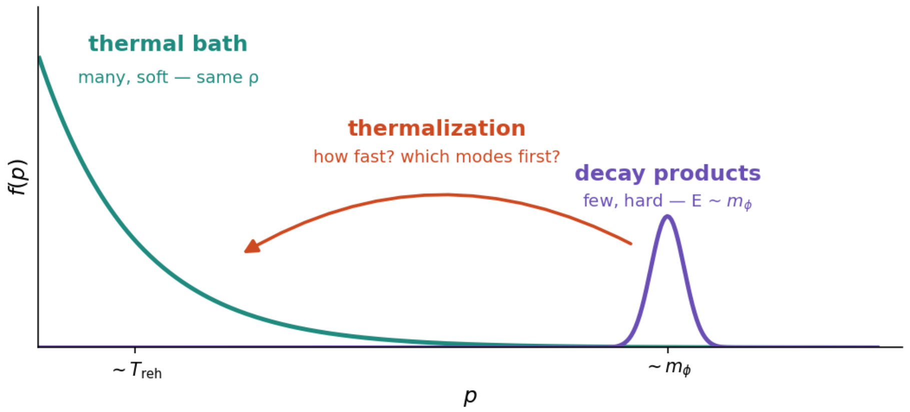

Thermalization (reaching full equilibrium) is governed by the Boltzmann equation with number-changing processes.

Dark matter genesis

The second episode on the timeline: a relic abundance, fixed by how χ falls out of equilibrium.

Where the standard machinery breaks:

- **Cannibal / SIMP** — relic set by $3 \to 2$ self-interactions
- **Freeze-in → freeze-out transition** — $\langle \sigma v \rangle$ averaging invalid
- **Dark-sector internal thermalization** — T_χ must emerge
- **Sub-threshold annihilation** — only particles above threshold can annihilate (→ Part II)

Different settings, one common demand: the **shape of f(p, t)**.

What existing tools assume

Every existing approach assumes the shape of f — the only question is how much:

Level	Solves for	Assumes	Tools
nBE	n(t)	equilibrium shape of f ; ⟨σv⟩	micrOMEGAs, MadDM, DarkSUSY, DRAKE
cBE	n(t), T(t)	kinetic equilibrium (shape up to T, μ)	DarkSUSY, DRAKE
relaxation time	f (p, t)	$C[f\,] \approx -(f - f_{eq})/\tau$: one timescale, no momentum-exchange information	—
fBE	f (p, t)	nothing about the shape	DRAKE (2 → 2) · BEST (any n → m)

The collision integral and its cost

$$\frac{\partial f_a}{\partial t} - Hp\frac{\partial f_a}{\partial p} = \sum_{\text{processes}} C_a[f] \qquad \frac{df_a(q,t)}{dt} = \sum_{\text{processes}} C_a[f]$$

$$C^{(k)}(\mathbf{p}) = \frac{1}{2E}\int d\Pi\, |\mathcal{M}|^2\, \Lambda \qquad \Lambda = \prod_{i\in\text{in}} f_i \prod_{j\in\text{out}} (1 \pm f_j) - \prod_{j\in\text{out}} f_j \prod_{i\in\text{in}} (1 \pm f_i) \qquad d\Pi = (2\pi)^4\delta^{(4)}\left(\sum_{i\in\text{in}} p_i - \sum_{j\in\text{out}} p_j\right)\prod_{l\neq k}\frac{d^3 p_l}{(2\pi)^3 2E_l}$$

- The right-hand side is an integral over every other particle — 3(n_total − 1) dimensions (2→2: 9, 2→3: 12) — with a δ^4 inside — at every grid point, every time step. No analytical reduction exists beyond 2 → 2.
- Can we just integrate it directly?

Handling the delta function

- In high dimensions there is one candidate: Monte Carlo, error ∝ $1/\sqrt{N\_eval}$. But the integrand contains δ^4: random samples never hit a zero-measure surface.

- Momentum (3 δ's): eliminated exactly. The conserved particle is determined, not sampled, and 3 dimensions are gone: $3(n\_total - 2)$ (2→2: 6, 2→3: 9). Roles: observed (pinned at p) · conserved · integrated.

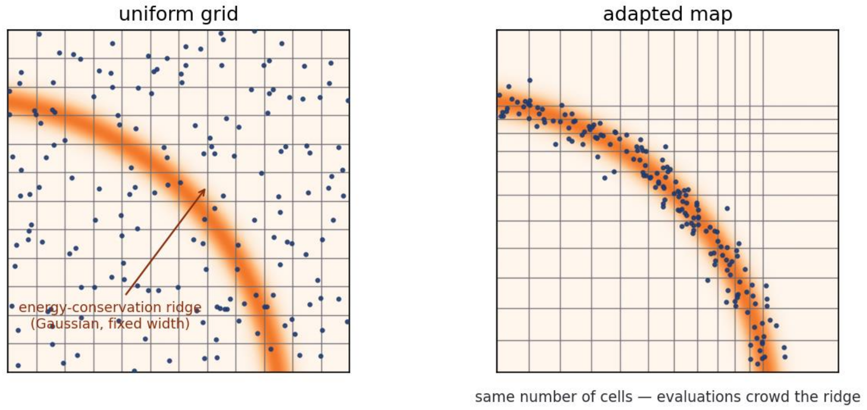

- Energy (1 δ): no variable left to solve for → narrow Gaussian

Vegas (Lepage '78; '21)

The identical-particle decomposition

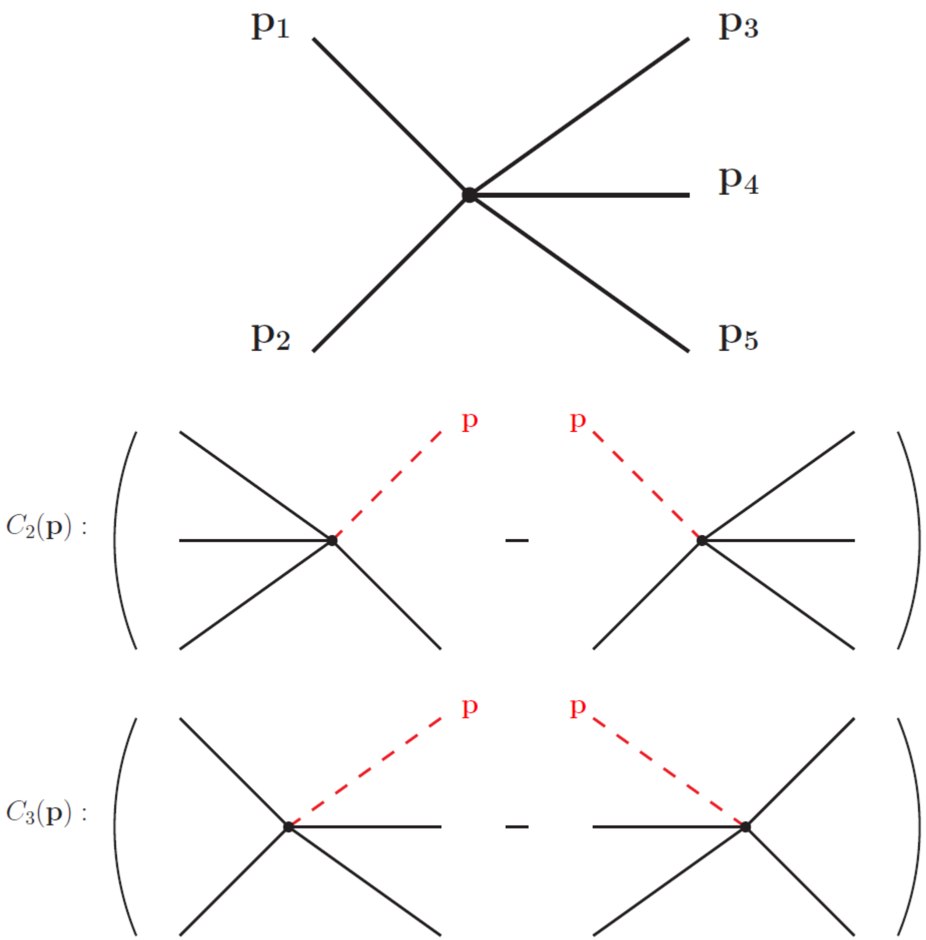

One particle is pinned — but which one, when they are identical?

For ϕϕ ↔ ϕϕϕ, p can sit on the 2-side or the 3-side

C2(p) — p on the 2-side:

gain: f3f4f5 (1+f (p))(1+f2)

loss: f (p)f2 (1+f3)(1+f4)(1+f5)

C3(p) — p on the 3-side:

gain: f1f2 (1+f (p))(1+f4)(1+f5)

loss: f (p)f4f5 (1+f1)(1+f2)

Two different integrals ⇒ keep both, slot-weighted: $C_\phi = 2 C_2 + 3 C_3$

Energy conservation requires the complete sum

Integrate E Ca(p) over all momenta using the full decomposition:

$$\int \frac{d^3p}{(2\pi)^3}\, E\, C_a(p) \propto \int d\Pi_{\text{all}} \left(\sum_{i\in\text{in}} E_i - \sum_{j\in\text{out}} E_j \right) |\mathcal{M}|^2\, \Lambda$$

vanishes by the energy–momentum δ-function — but only if the sum over all positions is complete.

- omit C2 for 2 ↔ 3: $\int E\, C_3\, p^2\, dp \neq 0 \Rightarrow$ systematic energy non-conservation (not an MC artifact)
- invisible in the nBE: momentum integration absorbs the sum over positions
- invisible to all previous fBE solvers: for 2 → 2, Cnα = Cnβ automatically

The rule, forced by energy conservation: C_a = n_α C_α + n_β C_β — every placement, weighted by its slots.

Validation: semi-analytical 2 → 2 benchmark

Independent cross-check, following Ala-Mattinen et al. ’22:

$$C_{\mathrm{BW}}(p_1) = \frac{2}{(2\pi)^4}\frac{1}{2E_1}\int dp_3\,\frac{p_3^2}{2E_3}\int dp_4\,\frac{p_4^2}{2E_4}\,F(p_1,p_3,p_4)\,f_3 f_4(1\pm f_1)(1\pm f_2)$$

- angular integrations done analytically; 9 → 2 dimensions
- energy conservation exact (E2 = E3 + E4 − E1, no Gaussian)
- ⇒ tests both the MC accuracy and the systematic error from the finite δ-width

Benchmark result

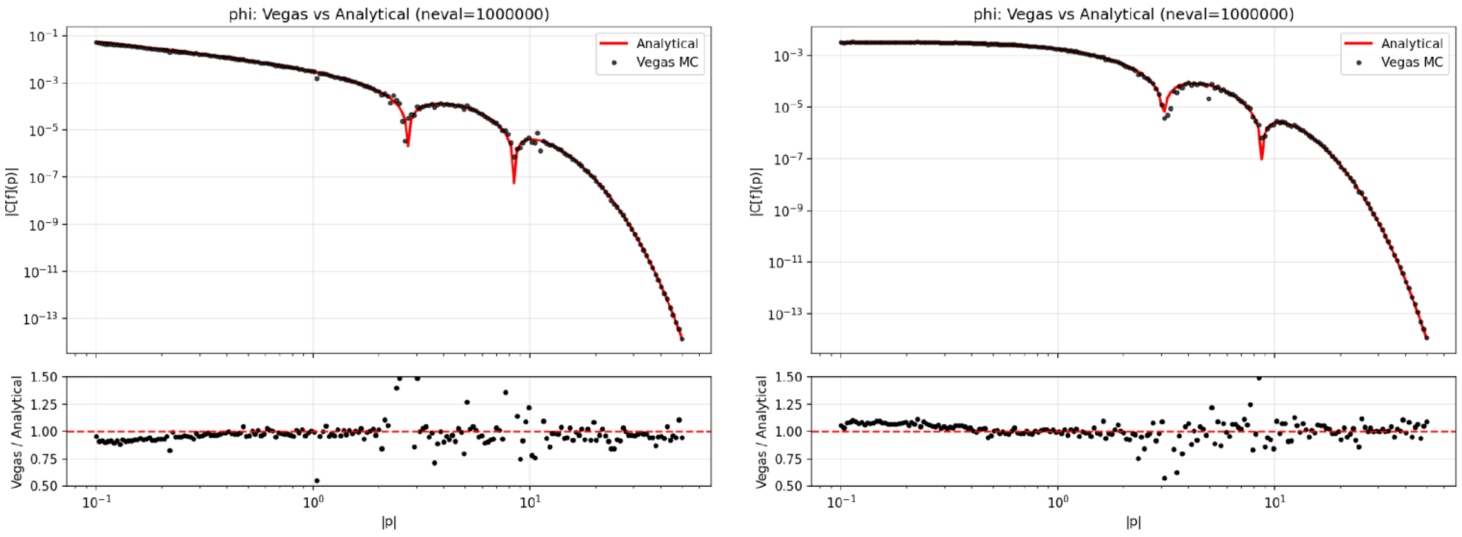

- What is compared: the collision integral itself, Cf — Vegas vs. the exact semi-analytical result
- Lower panels: the net rate (gain − loss) — ratio deviations near its zero crossings
- Agreement at the few-% level wherever the rate is sizable.

Thermalization: 2 → 2 elastic scattering

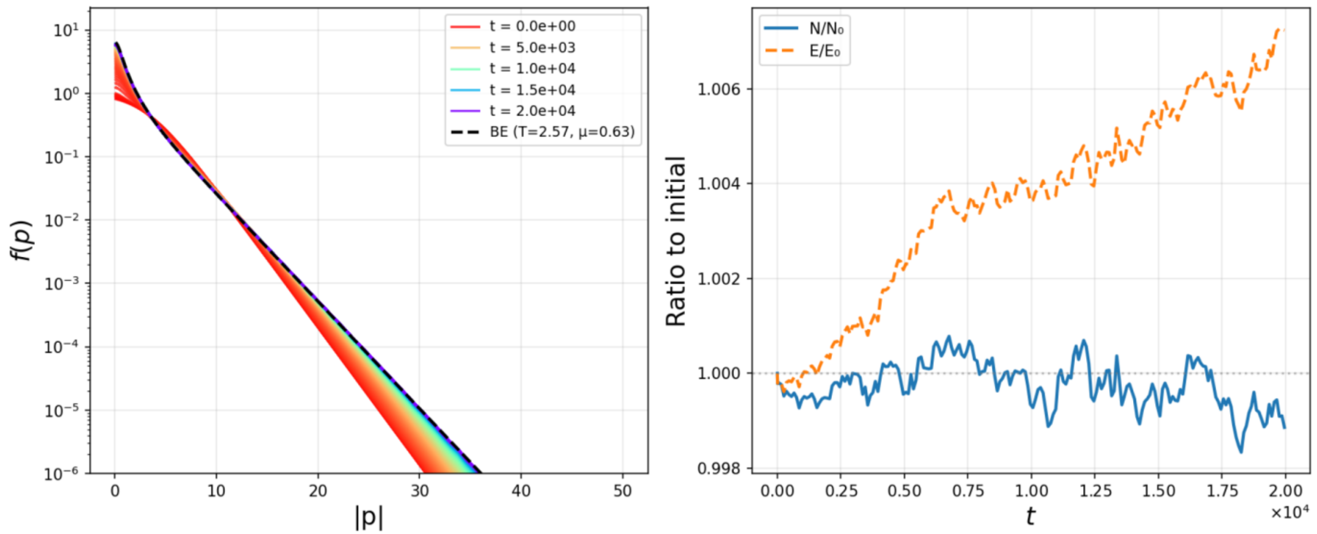

- What evolves: f(p) itself, from a non-thermal start (left) — while N and E are tracked as the check (right).
- **The endpoint is a prediction, not a fit** — (T_eq, μ) fixed by the conserved N and E.
- No thermal input anywhere — the Bose–Einstein form emerges; N, E conserved to ~1% (MC noise); rate balance emerges between independently adapted integrators

The key test: φφ ↔ φφφ

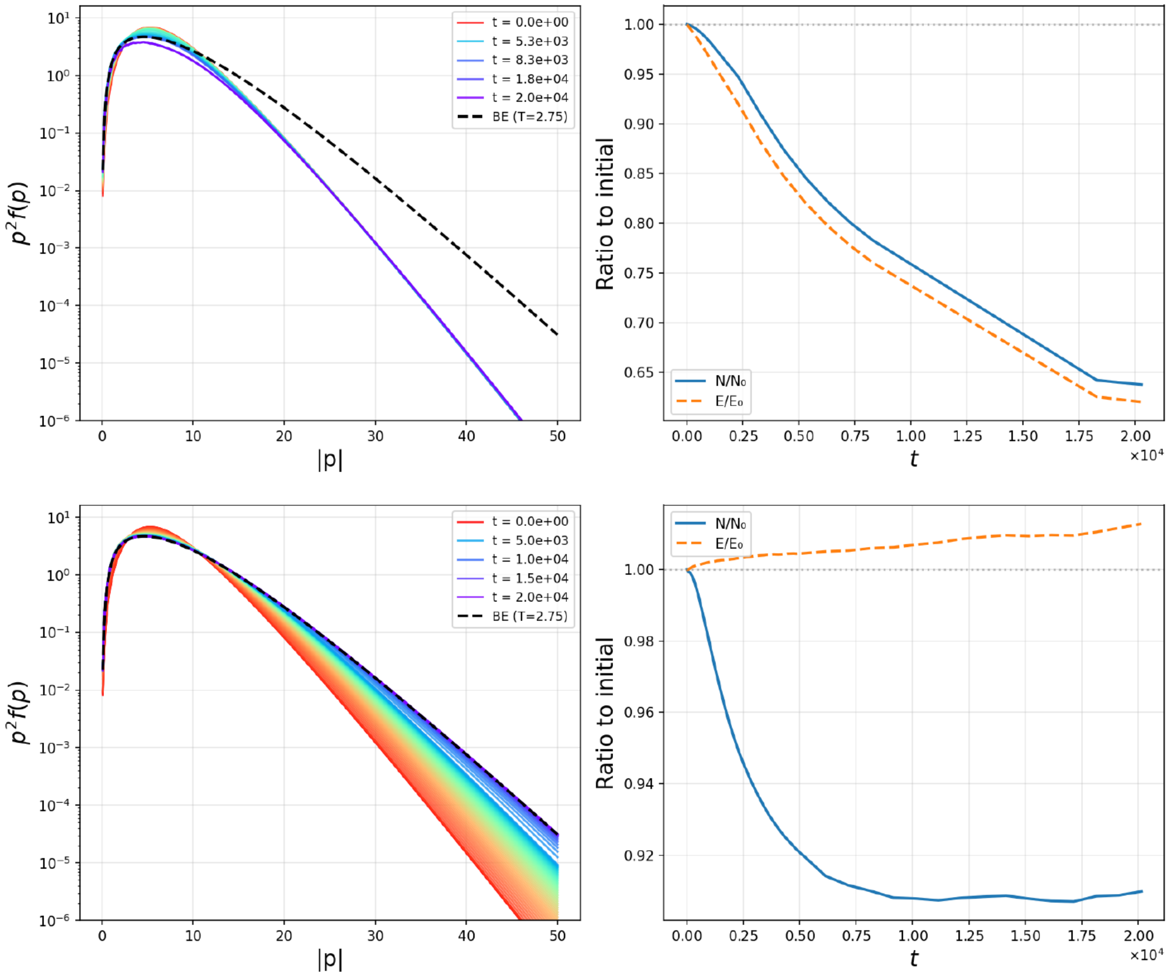

Top: naive implementation (C3 only)

- ~ 40% energy loss — systematic, not statistical
- wrong equilibrium: missing C2 distorts the production/absorption balance

Bottom: full 2 C2 + 3 C3

- energy conserved to ~ 1%
- f → Bose–Einstein with μ = 0 (chemical eq.: 3μ = 2μ)
- N/N0 → 0.9: the predicted ~10% decrease from the previous slide — net 3 → 2, μ → 0

From flat spacetime to the expanding universe

- Part I was a closed box: the system reaches equilibrium and stays. Now the bath cools underneath: **equilibrium itself becomes a moving target**

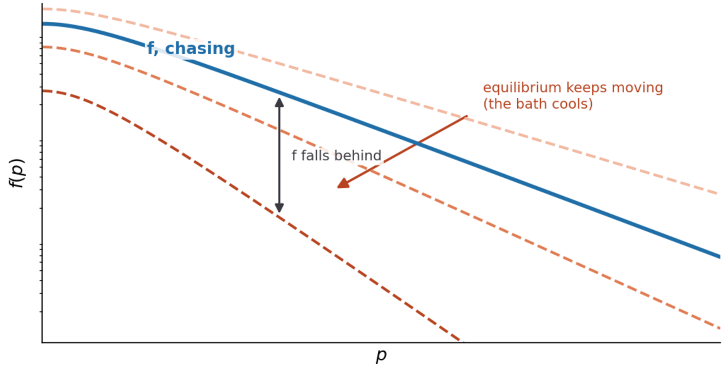

- The interactions chase that target while H sets the pace of the cooling — when the chase is lost, f falls behind: freeze-out

- Numerically this regime is **stiff**: steps with H, not Γ — exponential integrators

Dark matter freeze-out

- When annihilation falls behind the expansion, the abundance stops tracking equilibrium and freezes — this is the relic
- This story is usually told with one number, n(t) — but the thing that actually freezes is the full distribution f(p, t)
- The difference matters when annihilation depends on where a particle sits in f

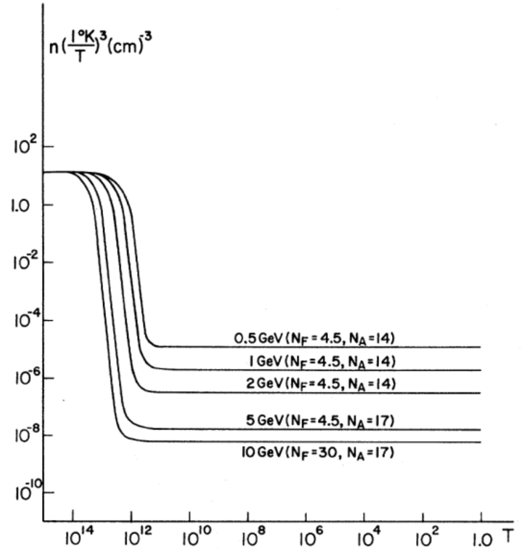

FIG. 1. n/T^3 vs T for a variety of special cases of m_L, N_F, and N_A.

Lee, Weinberg (1977)

Sub-threshold annihilation

$$\mathcal{L} \supset -\frac{\lambda}{4}\phi_1^2\phi_2^2 + y_f \phi_2 \bar{f} f$$

(annihilation & elastic) (keeps ϕ_2 thermal)

ϕ_1 = dark matter · ϕ_2 = bath, $m_2/m_1 = 1.1$ → sub-threshold

- As T drops, the typical ϕ_1 no longer has the energy to annihilate — only a window of f above threshold can
- Annihilation drains that window; scattering must refill it — the relic is set by that competition: by the shape of f

Dark matter relic abundance beyond kinetic equilibrium,
Binder, Bringmann, Gustafsson, Hryczuk (2021)

Sub-threshold annihilation: results

- Cross-checked against DRAKE

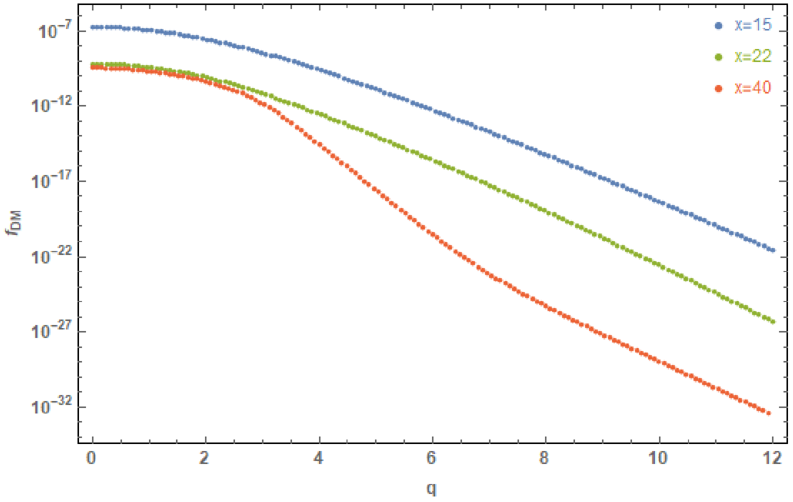

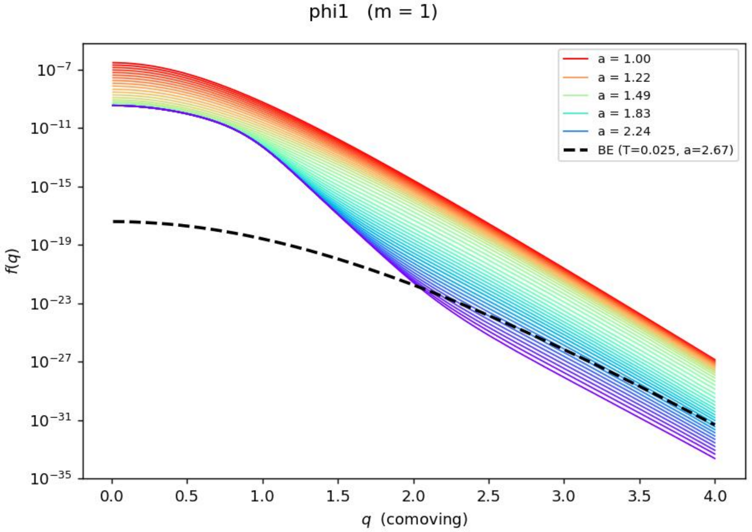

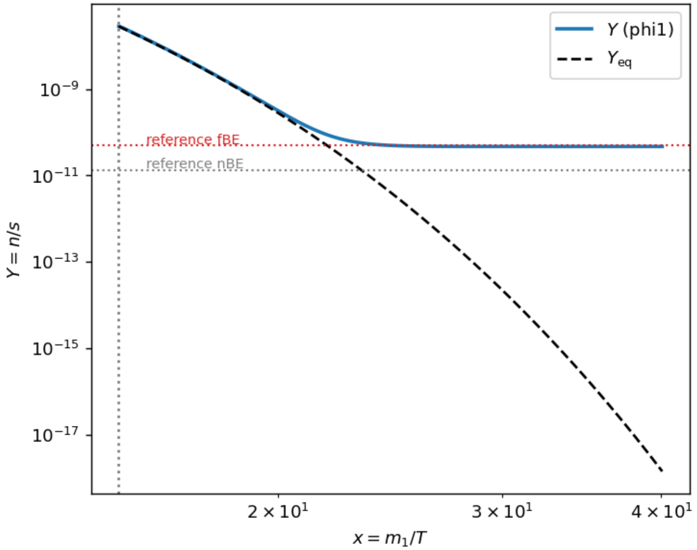

- Only particles above threshold can annihilate — the surviving low-energy bulk becomes the relic.
- Computed from the full f, that relic lands on the reference fBE line; the nBE misses it by ~×4.

Summary

So far — the full Boltzmann equation, solved directly with Monte Carlo:

- any $n \to m$ — the identical-particle decomposition, $C = n_\alpha C_\alpha + n_\beta C_\beta$
- validated against an exact $2 \to 2$ benchmark; $2 \leftrightarrow 3$ thermalization at %-level conservation
- through freeze-out: full-f relic on the reference-fBE line (preliminary)

Ahead — realistic reheating, dark matter models

github.com/best-hep/best   All the BEST.

The code in one slide

```
from besthep import BEST

solver = BEST(q_min=0.1, q_max=20.0,
              n_grid=64)
solver.initialize_species(
    'phi', init_f,
    stat='boson', mass=1.0)
solver.add_process(
    'cannibal',
    ['phi','phi'],
    ['phi','phi','phi'],
    matrix_element,
    coupling=1.0, neval=int(1e7))
for step in range(n_steps):
    solver.evolve_step(dt=dt)
```

- Python: numpy (vectorized batches), mpi4py, scipy
- dimensionality, phase-space construction, and the 2C2+3C3 decomposition: all automatic from the process declaration
- quantum statistics per species (boson/fermion/MB)
- time-dependent masses: set mass func('phi', ...)
- checkpoint/restart: single pickle file = full state (grids, f , Vegas maps, history)